

IN THE CLAIMS:

Please substitute the following claims for the same-numbered claims in the application:

1-28. (Cancelled).

29. (Currently Amended) A method for ~~dynamically developing~~ automatically providing a marketing strategy to address at least one specified merchant objective, the objective corresponding to a specified time period and a specified budget, the strategy being implemented across at least one marketing channel, the strategy including at least one initiative, the method comprising ~~the steps of:~~

inputting, by a merchant, said at least one specified merchant objective, said at least one specified merchant objective including said specified time period and said specified budget as constraints;

[[a.]] generating a plurality of possible marketing strategies, each of said plurality of possible marketing strategies comprising a set of initiatives, which are deployed together in a given sequence for said specified time period,

wherein said initiatives include any of bundling of products, cross-sells, up-sells, coupons, discounts, promotions, advertisements, surveys, and customer feedback;

[[b.]] determining an optimal marketing strategy ~~based on a state of a customer and constraints corresponding to marketing channels, such that said state of a customer comprises a purchase frequency and a monetary value of purchases,~~ from said plurality of possible marketing strategies, each of said plurality of possible marketing strategies corresponding to a policy comprising a sequence of various actions taken at different states encountered during said specified time period,

wherein each of said different states corresponds to a set of variables including any of customer profile, purchase frequency, and monetary value of purchase, associated with a customer at a time in said specified time period,

wherein said various actions include at least one randomized action;

wherein said various actions are constrained by a choice of a marketing channel,
wherein said various actions are deployed across said different states encountered
during said specified time period to provide a plurality of policies,

wherein each of said plurality of policies is evaluated in a context of a
reinforcement learning algorithm, in which values of each of said plurality of policies
corresponds to a vector of total expected rewards, said total expected rewards comprising a sum
of rewards corresponding to a monetary value for each of said various actions deployed across
each of said different states during said specified time period,

wherein said determining of said optimal marketing strategy comprises
determining an optimal policy for each state during said specified time period based on past data,
wherein said determining of said optimal policy comprises:

~~identifying a deterministic policy,~~
~~initializing a value of all possible states for said deterministic policy,~~
~~computing the value of a state for said deterministic policy, wherein said value~~
~~comprises a total expected reward for said state,~~
~~repeating said step of computing for all possible states,~~
~~constructing a new improved policy,~~
~~iteratively performing said steps of computing, repeating, and constructing until a~~
~~new improved policy remains unchanged for two subsequent iterations, and~~
~~selecting a policy with maximum value for the state as the optimal policy for the~~
~~given state;~~

evaluating, in order, said each state for said specified time period across
each of said plurality of policies, corresponding to said sum of rewards; and

identifying said optimal policy, associated with said optimal marketing
strategy, by a maximal value representing a maximal total expected reward for said optimal
policy; and

~~e. deploying~~ outputting, to said merchant, the determined optimal
marketing strategy;

~~d. recording customer response to the deployed optimal marketing strategy;~~

~~e. updating information corresponding to the state of a customer based on the recorded customer response; and~~

~~f. repeating steps b to e for the specified time period.~~

30. (Currently Amended) The method according to claim 29, ~~all the limitations of which are incorporated herein by reference~~, wherein the ~~step of~~ generating a plurality of possible marketing strategies comprises ~~the steps of~~:

selecting at least one initiative that enables an addressing of the at least one specified merchant objective;

determining sequences in which selected initiatives are deployed, if more than one initiative is selected; and

combining the selected initiatives in the determined sequences to obtain the plurality of possible marketing strategies.

31. (Currently Amended) The method according to claim 30, ~~all the limitations of which are incorporated herein by reference~~, further comprising varying parameters of initiatives to generate new initiatives, corresponding to new states during said specified time period.

32. (Currently Amended) The method according to claim 30, ~~all the limitations of which are incorporated herein by reference~~, further comprising varying deployment time of said initiatives.

33. (Currently Amended) The method according to claim 29, ~~all the limitations of which are incorporated herein by reference~~, wherein the ~~step of~~ determining an optimal marketing strategy further comprises ~~the steps of~~, after said determining of said optimal policy for each state based on past data:

identifying ~~[[the]]~~ a state of a customer, ~~[[the]]~~ a customer visiting a merchant, or ~~[[the]]~~ a customer being selected from a database of customers; and

identifying an optimal marketing strategy using the state of the customer, the identified optimal policy, and constraints corresponding to marketing channels.

34-35. (Cancelled).

36. (Currently Amended) The method according to claim ~~29~~ 35, ~~all the limitations of which are incorporated herein by reference~~, wherein ~~the step of computing the value of a state for the policy~~ evaluating each of said plurality of policies comprises ~~the steps of~~:

- computing transition probabilities from a given state to another state ~~for the policy~~;
- computing a value of expected immediate reward ~~for the policy in the~~ corresponding to said given state;
- computing a discounted expected value of a resulting state ~~for the policy~~; and
- computing a sum of expected immediate ~~reward~~ rewards and ~~[[the]]~~ a total discounted expected value corresponding to a sum of states.

37. (Canceled).

38. (Currently Amended) The method according to claim 33, ~~all the limitations of which are incorporated herein by reference~~, wherein ~~the step of~~ identifying an optimal marketing strategy comprises ~~the steps of~~:

- identifying the optimal policy for an identified customer state;
- ~~modeling~~ assigning customer's preferences for marketing channels, cost, and effectiveness of different marketing channels, and the specified budget as ~~effective~~ constraints;
- determining an optimal feasible policy based on the identified optimal policy and ~~effective~~ constraints corresponding to marketing channels; and
- determining ~~[[the]]~~ an optimal feasible marketing strategy from the optimal feasible policy.

39. (Currently Amended) The method according to claim 38, ~~all the limitations of which are incorporated herein by reference~~, wherein ~~the step of~~ determining an optimal feasible policy based on ~~effective~~ constraints corresponding to marketing channels comprises mapping the

optimal policy uniquely to a closest feasible optimal policy based on the ~~effective~~ constraints, if the ~~effective~~ constraints are not satisfied by the optimal policy.

40-42. (Canceled).

43. (Currently Amended) A system for ~~dynamically developing~~ automatically providing a marketing strategy to address at least one specified merchant objective, the objective corresponding to a specified time period and a specified budget, the strategy being implemented across at least one marketing channel, the strategy including at least one initiative, the system comprising:

a memory for storing said at least one specified merchant objective, which is inputted by a merchant, said at least one specified merchant objective including said specified time period and said specified budget as constraints; and

a microprocessor configured to:

~~a generator operable for generating~~ generate a plurality of possible marketing strategies, each of said plurality of possible marketing strategies comprising a set of initiatives, which are deployed together in a given sequence for said specified time period,

wherein said initiatives include any of bundling of products, cross-sells, up-sells, coupons, discounts, promotions, advertisements, surveys, and customer feedback;

~~a first unit operable for determining~~ determine an optimal marketing strategy ~~based on state of a customer and constraints corresponding to marketing channels, such that said state of a customer comprises a purchase frequency and a monetary value of purchases, from said~~ plurality of possible marketing strategies, each of said plurality of possible marketing strategies corresponding to a policy comprising a sequence of various actions taken at different states encountered during said specified time period,

wherein each of said different states corresponds to a set of variables including any of customer profile, purchase frequency, and monetary value of purchase, associated with a customer at a time in said specified time period,

wherein said various actions include at least one randomized action;

wherein said various actions are constrained by a choice of a marketing channel and said specified budget,

wherein said various actions are deployed across said different states encountered during said specified time period to provide a plurality of policies,

wherein each of said plurality of policies is evaluated in a context of a reinforcement learning algorithm, in which values of each of said plurality of policies corresponds to a vector of total expected rewards, said total expected rewards comprising a sum of rewards corresponding to a monetary value for each of said various actions deployed across each of said different states during said specified time period,

wherein said determining of said optimal marketing strategy comprises determining an optimal policy for each state during said specified time period based on past data, wherein said determining of said optimal policy comprises:

~~identifying a deterministic policy,~~

~~initializing a value of all possible states for said deterministic policy,~~

~~computing the value of a state for said deterministic policy, wherein said value comprises a total expected reward for said state,~~

~~repeating said step of computing for all possible states,~~

~~constructing a new improved policy,~~

~~iteratively performing said steps of computing, repeating, and constructing until a new improved policy remains unchanged for two subsequent iterations, and~~

~~selecting a policy with maximum value for the state as the optimal policy for the given state;~~

~~a first sub-unit operable for determining all possible states of a second unit operable for deploying the determined optimal marketing strategy;~~

~~a recorder operable for recording customer response to the deployed optimal marketing strategy; and~~

~~a third unit operable for updating information corresponding to the state of a customer based on the recorded customer response~~

evaluating, in order, said each state for said specified time period across each of said plurality of policies, corresponding to said sum of rewards; and
identifying said optimal policy, associated with said optimal marketing strategy, by a maximal value representing a maximal total expected reward for said optimal policy; and
outputting, to said merchant, the optimal marketing strategy.

44. (Currently Amended) The system according to claim 43, ~~all the limitations of which are incorporated herein by reference~~, wherein said ~~generator~~ generating a plurality of possible marketing strategies comprises:

~~a selector operable for~~ selecting at least one initiative that enables an addressing of the at least one specified merchant objective;

~~a first sub-unit operable for~~ determining sequences in which selected initiatives are deployed, when more than one initiative is selected; and

~~a second sub-unit~~ for combining the selected initiatives in the determined sequences to obtain the plurality of possible marketing strategies.

45. (Currently Amended) The system according to claim 43, ~~all the limitations of which are incorporated herein by reference~~, wherein the ~~first unit~~ identifying an optimal marketing strategy comprises:

~~a second sub-unit operable for determining an optimal policy for each state based on past data;~~

~~a third sub-unit operable for~~ identifying ~~[[the]]~~ a state of a customer, ~~[[the]]~~ a customer visiting a merchant, or ~~[[the]]~~ a customer being selected from a database of customers;

~~a fourth sub-unit operable for~~ identifying ~~[[the]]~~ an optimal policy for an identified customer state;

~~a fifth sub-unit operable for modeling~~ assigning customer's preferences for marketing channels, cost, and effectiveness of different marketing channels, and the specified budget as effective constraints;

~~a sixth sub-unit operable for determining an optimal feasible policy based on effective constraints corresponding to marketing channels; and~~

~~a seventh sub-unit operable for determining [[the]] an optimal feasible marketing strategy from the optimal feasible policy.~~

46. (Cancelled).

47. (Currently Amended) The system according to claim ~~[[46]]~~ 45, ~~all the limitations of which are incorporated herein by reference~~, wherein the ~~fourth component comprises a selector operable for selecting the marketing strategy that maximizes a value for the state over all marketing strategies for a given state~~ determining an optimal feasible policy based on constraints corresponding to marketing channels comprises mapping the optimal policy uniquely to a closest feasible optimal policy based on the constraints, if the constraints are not satisfied by the optimal policy.

48. (Canceled).

49. (Currently Amended) A program storage device readable by computer, tangibly embodying a program of instructions executable by the computer to perform a method for ~~dynamically developing~~ automatically providing a marketing strategy to address at least one specified merchant objective, the objective corresponding to a specified time period and a specified budget, the strategy being implemented across at least one marketing channel, the strategy including at least one initiative, the method comprising:

inputting, by a merchant, said at least one specified merchant objective, said at least one specified merchant objective including said specified time period and said specified budget as constraints;

generating a plurality of possible marketing strategies, each of said plurality of possible marketing strategies comprising a set of initiatives, which are deployed together in a given sequence for said specified time period,

wherein said initiatives include any of bundling of products, cross-sells, up-sells, coupons, discounts, promotions, advertisements, surveys, and customer feedback;

determining an optimal marketing strategy based on a state of a customer and constraints corresponding to marketing channels, such that said state of a customer comprises a purchase frequency and a monetary value of purchases, from said plurality of possible marketing strategies, each of said plurality of possible marketing strategies corresponding to a policy comprising a sequence of various actions taken at different states encountered during said specified time period,

wherein each of said different states corresponds to a set of variables including any of customer profile, purchase frequency, and monetary value of purchase, associated with a customer at a time in said specified time period,

wherein said various actions include at least one randomized action;

wherein said various actions are constrained by a choice of a marketing channel and said specified budget,

wherein said various actions are deployed across said different states encountered during said specified time period to provide a plurality of policies,

wherein each of said plurality of policies is evaluated in a context of a reinforcement learning algorithm, in which values of each of said plurality of policies corresponds to a vector of total expected rewards, said total expected rewards comprising a sum of rewards corresponding to a monetary value for each of said various actions deployed across each of said different states during said specified time period,

wherein said determining of said optimal marketing strategy comprises determining an optimal policy for each state during said specified time period based on past data, wherein said determining of said optimal policy comprises:

~~identifying a deterministic policy,~~

~~initializing a value of all possible states for said deterministic policy,~~

~~computing the value of a state for said deterministic policy, wherein said value comprises a total expected reward for said state,~~

~~repeating said step of computing for all possible states,~~

~~constructing a new improved policy;~~
~~iteratively performing said steps of computing, repeating, and constructing until a~~
~~new improved policy remains unchanged for two subsequent iterations, and~~
~~selecting a policy with maximum value for the state as the optimal policy for the~~
~~given state;~~
evaluating, in order, said each state for said specified time period across
each of said plurality of policies, corresponding to said sum of rewards; and
identifying said optimal policy, associated with said optimal marketing
strategy, by a maximal value representing a maximal total expected reward for said optimal
policy; and
~~deploying~~ outputting, to said merchant, the determined optimal marketing
strategy;
~~recording customer response to the deployed optimal marketing strategy; and~~
~~updating information corresponding to the state of a customer based on the recorded~~
~~customer response.~~

50. (Currently Amended) The program storage device according to claim 49, ~~all the~~
~~limitations of which are incorporated herein by reference~~, wherein the ~~step of~~ generating a
plurality of possible marketing strategies comprises:

selecting at least one initiative that enables an addressing of the at least one specified
merchant objective;

determining sequences in which selected initiatives are deployed, when more than one
initiative is selected; and

combining the selected initiatives in the determined sequences to obtain the plurality of
possible marketing strategies.

51. (Currently Amended) The program storage device according to claim 49, ~~all the~~
~~limitations of which are incorporated herein by reference~~, wherein the ~~step of~~ determining an
optimal marketing strategy comprises:

~~determining an optimal policy for each state based on past data;~~
identifying ~~[[the]]~~ a state of a customer, ~~[[the]]~~ a customer visiting a merchant, or ~~[[the]]~~ a customer being selected from a database of customers;
identifying ~~[[the]]~~ an optimal policy for an identified customer state;
~~modeling~~ assigning customer's preferences for marketing channels, cost, and effectiveness of different marketing channels, and the specified budget as ~~effective~~ constraints;
determining an optimal feasible policy based on ~~effective~~ constraints corresponding to marketing channels; and
determining ~~[[the]]~~ an optimal feasible marketing strategy from the optimal feasible policy.

52-55. (Cancelled).

56. (Currently Amended) A method for ~~dynamically developing~~ automatically providing a marketing strategy to address at least one specified merchant objective, the objective corresponding to a specified time period and a specified budget, the strategy being implemented across at least one marketing channel, the strategy including at least one initiative, the method comprising ~~the steps of:~~

- ~~a. generating a plurality of marketing strategies;~~
- ~~b. determining all possible states of customers, such that said states of customers comprise a purchase frequency and a monetary value of purchases;~~
- ~~c. determining an optimal policy for each state based on past data;~~
- ~~d. identifying the state of a customer, the customer visiting a merchant or the customer being selected from a database of customers;~~
- ~~e. identifying the optimal policy for an identified customer state;~~
- ~~f. modeling customer's preferences for marketing channels, cost and effectiveness of different marketing channels, and the specified budget as effective constraints;~~
- ~~g. determining an optimal feasible policy based on the identified optimal policy and effective constraints corresponding to marketing channels;~~

~~h. determining an optimal marketing strategy from the optimal feasible policy, wherein said determining of said optimal marketing strategy comprises determining an optimal policy for each state based on past data, wherein said determining of said optimal policy comprises:~~

~~identifying a deterministic policy,~~

~~initializing a value of all possible states for said deterministic policy,~~

~~computing the value of a state for said deterministic policy, wherein said value comprises a total expected reward for said state,~~

~~repeating said step of computing for all possible states,~~

~~constructing a new improved policy,~~

~~iteratively performing said steps of computing, repeating, and constructing until a new improved policy remains unchanged for two subsequent iterations, and~~

~~selecting a policy with maximum value for the state as the optimal policy for the given state;~~

~~i. deploying the determined optimal marketing strategy;~~

~~j. recording customer response to the deployed marketing strategy;~~

~~k. identifying a resulting state of the customer;~~

~~l. updating values of the state of the customer;~~

~~m. updating the optimal policy; and~~

~~n. repeating steps c to m for the specified time period~~

inputting, by a merchant, said at least one specified merchant objective, said at least one specified merchant objective including said specified time period and said specified budget as constraints;

generating a plurality of possible marketing strategies, each of said plurality of possible marketing strategies comprising a set of initiatives, which are deployed together in a given sequence for said specified time period,

wherein said initiatives include any of bundling of products, cross-sells, up-sells, coupons, discounts, promotions, advertisements, surveys, and customer feedback;

determining an optimal marketing strategy from said plurality of possible marketing strategies, each of said plurality of possible marketing strategies corresponding to a policy

comprising a sequence of various actions taken at different states encountered during said specified time period,

wherein each of said different states corresponds to a set of variables including any of customer profile, purchase frequency, and monetary value of purchase, associated with a customer at a time in said specified time period,

wherein said various actions include at least one randomized action;

wherein said various actions are constrained by a choice of a marketing channel and said specified budget,

wherein said various actions are deployed across said different states encountered during said specified time period to provide a plurality of policies,

wherein each of said plurality of policies is evaluated in a context of a reinforcement learning algorithm, in which values of each of said plurality of policies corresponds to a vector of total expected rewards, said total expected rewards comprising a sum of rewards corresponding to a monetary value for each of said various actions deployed across each of said different states during said specified time period,

wherein said determining of said optimal marketing strategy comprises determining an optimal policy for each state during said specified time period based on past data, wherein said determining of said optimal policy comprises:

evaluating, in order, said each state for said specified time period across each of said plurality of policies, corresponding to said sum of rewards; and

identifying said optimal policy, associated with said optimal marketing strategy, by a maximal value representing a maximal total expected reward for said optimal policy;

outputting, to said merchant, the optimal marketing strategy;

recording customer response to the outputted optimal marketing strategy; and

updating information corresponding to a state of a customer based on the recorded customer response.